

IBM Docket No. AUS920010967US1

1

TITLE OF THE INVENTION

Employing Speech Recognition and Key Words to Improve Customer Service

5 CROSS-REFERENCES TO RELATED APPLICATIONS

The present application is related to a co-pending application entitled Employing Speech Recognition and Capturing Customer Speech to Improve Customer Service, filed on even date herewith, assigned to the assignee of the present application, and herein 10 incorporated by reference.

FIELD OF THE INVENTION

The present invention relates generally to information handling, and more particularly to methods and systems employing 15 computerized speech recognition and key words to improve customer service.

BACKGROUND OF THE INVENTION

Many approaches to speech transmission and speech recognition 20 have been proposed in the past, including the following examples: U.S. Pat. No. 6,100,882 (Sharman, et al., Aug. 8, 2000), "Textual Recording of Contributions to Audio Conference Using Speech Recognition," relates to producing a set of minutes for a teleconference. U.S. Pat. No. 6,243,454 (Eslambolchi, June 5, 25 2001), "Network-Based Caller Speech Muting," relates to a method for muting a caller's outgoing speech to defeat transmission of ambient noise, as with a caller in an airport. U.S. Pat. No. 5,832,063 (Vysotsky et al., Nov. 3, 1998), relates to speaker-independent recognition of commands, in parallel with speaker-dependent recognition of names, words or phrases, for speech - 30

activated telephone service. However, the above-mentioned examples address substantially different problems (i.e. problems of telecommunications service), and thus are significantly different from the present invention.

5

There are methods and systems in use today that utilize automatic speech recognition to replace human customer service representatives. Automatic speech recognition systems are capable of performing some tasks; however, a customer may need or prefer to actually speak with another person in many cases. Thus there is a need for systems and methods that use both automatic speech recognition, and human customer service representatives, automatically utilizing key words in customer speech to improve the customer service rendered by humans.

15

SUMMARY OF THE INVENTION

The present invention comprises capturing a customer's speech, recognizing a key word in the customer's speech, searching a database, and retrieving information from the database. The retrieving is a real-time process, completed during a conversation involving the customer and a customer service representative.

25

Consider some examples. It would be useful to identify key words spoken by a customer, through speech recognition technology. There would be no need to prolong a transaction by having a customer service representative manually input key words, or manually direct a search. Consider using a key word to automatically search a database and retrieve information from the database. The retrieved information could be automatically

30

provided to a customer service representative. Another example would be automatically initiating action, based on said key word, to solve a problem affecting said customer. This may involve preparing an e-mail message, preparing an order form, preparing an address label, routing a telephone call, or other similar tasks. These examples could apply to any business or organization providing goods or services.

BRIEF DESCRIPTION OF THE DRAWINGS

A better understanding of the present invention can be obtained when the following detailed description is considered in conjunction with the following drawings. The use of the same reference symbols in different drawings indicates similar or identical items.

FIG. 1 illustrates a simplified example of a computer system capable of performing the present invention.

FIG. 2 is a high- level block diagram illustrating an example of a system employing computerized speech recognition and key words, according to the teachings of the present invention.

FIG. 3 illustrates selected operations of another exemplary system, employing computerized speech recognition and key words.

FIG. 4 is a flow chart illustrating an example of a speech recognition and analysis function, according to the teachings of the present invention.

FIG. 5 is a flow chart illustrating an example of a research assistant function.

FIG. 6 is a flow chart illustrating an example of a resolution assistant function.

FIG. 7 illustrates selected operations of another exemplary system, employing a research assistant function or a resolution

assistant function.

DETAILED DESCRIPTION

The examples that follow involve the use of one or more computers and may involve the use of one or more communications networks. The present invention is not limited as to the type of computer on which it runs, and not limited as to the type of network used.

As background information for the present invention, reference is made to the book by M. R. Schroeder, Computer Speech:

Recognition, Compression, Synthesis, 1999, Springer-Verlag, Berlin, Germany. This book provides an overview of speech technology, including automatic speech recognition and speaker identification. This book provides introductions to two common types of speech recognition technology: statistical hidden Markov modeling, and neural networks. Reference is made to the book edited by Keith Ponting, Computational Models of Speech Pattern Processing, 1999, Springer-Verlag, Berlin, Germany. This book contains two articles that are especially useful as background information for the present invention. First, the article by Steve Young, "Acoustic Modeling for Large Vocabulary Continuous Speech Recognition," at pages 18 - 39, provides a description of benchmark tests for technologies that perform speaker-independent recognition of continuous speech. (At the time of that publication, the state-of-the-art performance on "clean speech dictation within a limited domain such as business news" was around 7% word error [WER].) Secondly, the article by Jean-Paul Haton, "Connectionist and Hybrid Models for Automatic Speech Recognition," pages 54 - 66, provides a survey of research on hidden Markov modeling and neural networks.

The following are some examples of speech recognition technology that would be suitable for implementing the present invention. Large-vocabulary technology is available from IBM in the VIAVOICE and WEBSPHERE product families. SPHINX speech-recognition technology is freely available via the World Wide Web as open source software, from the Computer Science Division of Carnegie Mellon University, Pittsburgh, Pennsylvania. SPHINX 2 is described as real-time, large-vocabulary, and speaker-independent. SPHINX 3 is slower but more accurate, and may be suitable for transcription for example. Other technology similar to the above-mentioned examples also may be used.

Another technology that may be suitable for implementing the present invention is extensible markup language (XML), and in particular, VoiceXML. XML provides a way of containing and managing information that is designed to handle data exchange among various data systems. Thus it is well-suited to implementation of the present invention. Reference is made to the book by Elliotte Rusty Harold and W. Scott Means, XML in a Nutshell (O'Reilly & Associates, 2001). As a general rule XML messages use "attributes" to contain information about data, and "elements" to contain the actual data. As background information for the present invention, reference is made to the article by Lee Anne Phillips, "VoiceXML and the Voice/ Web Environment: Visual Programming Tools for Telephone Application Development," Dr. Dobb's Journal, Vol. 26, Issue 10, pages 91 - 96, October 2001. One example described in the article is a currency-conversion application. It receives input, via speech and telephone, of an amount of money. It responds with an equivalent in another currency either via speech or via data display.

The following are definitions of terms used in the description of the present invention and in the claims:

"Computer-usable medium" means any carrier wave, signal or transmission facility for communication with computers, and any kind of computer memory, such as floppy disks, hard disks, Random Access Memory (RAM), Read Only Memory (ROM), CD-ROM, flash ROM, non-volatile ROM, and non-volatile memory.

"Customer" means any buyer, client, consumer, patient, patron, or user.

"Customer service" means interaction with a customer of any business, organization, or profession.

"Customer service representative" or "service representative" means any professional or other person who interacts with a customer, including an agent, assistant, broker, banker, consultant, engineer, legal professional, medical professional, or sales person.

"Key word" or "keyword" means any name, number, phrase, term, word or other characters that may be used in searching for information in a database.

"Storing" data or information, using a computer, means placing the data or information, for any length of time, in any kind of computer memory, such as floppy disks, hard disks, Random Access Memory (RAM), Read Only Memory (ROM), CD-ROM, flash ROM, non-volatile ROM, and non-volatile memory.

FIG. 1 illustrates a simplified example of an information handling system that may be used to practice the present invention. The invention may be implemented on a variety of hardware platforms, including personal computers, workstations, servers, and embedded systems. The computer system of FIG. 1 has

at least one processor 110. Processor 110 is interconnected via system bus 112 to random access memory (RAM) 116, read only memory (ROM) 114, and input/output (I/O) adapter 118 for connecting peripheral devices such as disk unit 120 and tape drive 140 to bus 112. The system has analog/digital converter 162 for connecting the system to telephone hardware 164 and public switched telephone network 160. The system has user interface adapter 122 for connecting keyboard 124, mouse 126, or other user interface devices such as audio output device 166 and audio input device 168 to bus 112. The system has communication adapter 134 for connecting the information handling system to a data processing network 150, and display adapter 136 for connecting bus 112 to display device 138. Communication adapter 134 may link the system depicted in FIG. 1 with hundreds or even thousands of similar systems, or other devices, such as remote printers, remote servers, or remote storage units. The system depicted in FIG. 1 may be linked to both local area networks (sometimes referred to as Intranets) and wide area networks, such as the Internet.

While the computer system described in FIG. 1 is capable of executing the processes described herein, this computer system is simply one example of a computer system. Those skilled in the art will appreciate that many other computer system designs are capable of performing the processes described herein.

FIG. 2 is a high- level block diagram illustrating an example of a system, 230, employing computerized speech recognition and capturing customer speech. System 230 is shown receiving speech input from two or more parties to a telephone conversation,

including a first speaker (such as customer service representative 220 for example). System 230 blocks a portion of the speech input that originates from the first speaker (service representative 220) and performs speech recognition on the

5 remaining portion of the speech input. The blocking and performing speech recognition are real-time processes, completed during a conversation. System 230 includes various components.

10 De-clutter component 231 de-clutters the speech input from service representatives 220 and 225 and customer 210 for better automatic processing, by removing all but the pertinent words spoken by the customer. This will be explained in more detail below.

15 After capturing customer 210's speech, system 230 recognizes a key word in customer 210's speech. Based on said key word, system 230 searches a database 260, and retrieves information from database 260. System 230 includes a speech recognition and analysis component 232, that may be implemented with well-known speech recognition technologies.

20 System 230 includes a key word database or catalog 235 that comprises a list of searchable terms. An example is a list of terms in a software help index. As indicated by the dashed line, key word database 235 may be incorporated into system 230, or may be independent of, but accessible to, system 230. Key word database 235 may be implemented with database management software such as ORACLE, SYBASE, or IBM's DB2, for example. An organization may create key word database 235 by pulling information from existing databases containing customer data and product data, for example. A customer name is an example of a key

word. A text extender function, such as that available with IBM's DB2, would allow a spoken name such as "Petersen" to be retrieved through searches of diverse spellings like "Peterson" or "Pedersen." Other technology similar to the above-mentioned examples also may be used.

System 230 may also include research assistant component 233, that would automate data-retrieval functions involved when service representatives 220 and 225 assist customer 210. Data may be retrieved from one or more databases 260, either directly or via network 250. Resolution assistant component 234 would automate actions to resolve problems for customer 210. Resolution assistant component 234 may employ mail function 240, representing an e-mail application, or conventional, physical mail or delivery services. Thus information, goods, or services could be supplied to customer 210.

In this example, service representatives 220 and 225 are shown interacting with customer 210 via telephone, represented by telephone hardware 211, 221, and 226. A similar system could be used for face-to-face interactions. Service representatives 220 and 225 are shown interacting with system 230 via computers 222 and 227. This represents a way to display information that is retrieved from database 260, to service representatives 220 and 225. Service representatives 220 and 225 may be located at the same place, or at different places.

FIG. 3 illustrates selected operations of another exemplary system, employing computerized speech recognition and capturing customer speech. Customer speech is symbolized by the letters in

5 bubble 310. A service representative's speech is symbolized by the letters in bubble 320. De-clutter component 231 is shown receiving speech input (arrows 315 and 325) from two speakers, including a first speaker (service representative 220); blocking a portion of the speech input that originates from the first speaker (service representative 220); and processing the remaining portion of the speech input with a computer (speech recognition and analysis component 232). The blocking and processing are real-time processes, completed during a 10 conversation. Speech recognition and analysis component 232 is shown receiving speech input (arrow 330) from a customer 210. Speech recognition and analysis component 232 performs speech recognition on the speech input to generate a text equivalent, and parses the text to identify key words (arrows 332 and 334). 15

20 15 The key words at arrows 332 and 334 ("patch," "floating point," and "compiler") are examples that may arise in the computer industry. Also consider an example from the financial services industry. A customer may ask for help regarding an Individual Retirement Account. A service representative may ask: "Did you say that you wanted help with a Roth IRA?" The customer may respond: "No, I need help with a standard rollover IRA." The present invention would block that portion of the speech input that originates from the service representative, and process the 25 remaining portion of the speech input that contains "rollover" and "IRA" as examples of key words.

30 Research assistant component 233 is shown searching for an occurrence of key words 334 in a database 360, retrieving information from database 360, and providing retrieved

information (arrow 345) to service representative 220. The retrieving is completed during a conversation involving customer 210 and service representative 220. Thus research assistant component 233 would automate data-retrieval functions involved when service representative 220 assists customer 210. Research assistant component 233 may be implemented with well-known search engine technologies. Databases shown at 360 may contain customer information, product information or problem management information, for example.

10

Resolution assistant component 234 is shown searching for an occurrence of a key word 332 in a database 260, retrieving information from database 260, and sending mail (arrow 340) to customer 210. Thus resolution assistant component 234 initiates action, based on a key word 332, to solve a problem affecting customer 210. Resolution assistant component 234 may initiate one or more tasks such as sending a message by e-mail, preparing an order form, preparing an address label, or routing a telephone call. Resolution assistant component 234 may be implemented with well-known search engine and e-mail technologies, for example. Databases shown at 260 may contain customer names and addresses, telephone call - routing information, problem management information, product update information, order forms, or advisory bulletins for example.

20

FIG. 4 is a flow chart illustrating an example of a speech recognition and analysis function, according to the teachings of the present invention. The function begins at block 410, receiving speech input from a customer. At block 420 the function performs speech recognition on said speech input to generate a

30

text equivalent. At block 430 the function parses said text to identify at least one key word. At block 440 the function searches for an occurrence of said key word in at least one database. At block 450 the function retrieves information from at least one database, and passes the information to a research assistant function or a resolution assistant function. The retrieving is completed during a conversation involving the customer and a customer service representative. In another example, search and retrieval could be performed by a research assistant function or a resolution assistant function. At block 460 the function creates or updates a problem management record (PMR) to document research or actions to solve a problem affecting the customer.

If the problem is not resolved at this point, the "No" branch is taken at decision 470, and the function returns to block 420. For example, this may represent continued conversation involving the customer and a customer service representative, and continued speech input. On the other hand, this may represent a decision to transfer handling to another customer service representative, who would continue the conversation with the customer. This would result in further speech input from the customer. The second representative may be more skilled in the problem area. The two representatives may be at the same location or at different locations. The two representatives may use the same system, or separate systems, to execute the process shown in FIG. 4. A decision to transfer handling to another customer service representative may be based on one or more key words (e.g. a customer name, product name, or problem description), automatically recognized in the customer's speech.

On the other hand, if the problem is resolved at this point, the "Yes" branch is taken at decision 470, and next the function updates one or more databases at block 480. For example, a new key word may be identified and added to a key word database for future use, or information about a problem and its solution may be added to a customer database, or a product database. The function terminates at block 490 (e.g. at the end of a conversation).

FIG. 5 is a flow chart illustrating an example of a research assistant function. The function begins at block 510, which may represent the research assistant function being called by another function such as the speech recognition and analysis function illustrated in FIG. 4. At block 520, the function receives search results, i.e. information retrieved from one or more databases. In another example, search and retrieval could be performed by a research assistant function. At block 530, the function ranks the search results, to prepare the information for use by a customer service representative. For example, information predicted to be more useful could be highlighted, listed first, or otherwise presented more prominently than other information. At block 540, the function provides the retrieved information to a customer service representative (e.g. via audio output, or by displaying the information on a screen, or printing information on paper). The function terminates at block 550 (e.g. at the end of a customer service representative's interaction with a customer).

FIG. 6 is a flow chart illustrating an example of a resolution assistant function. The function begins at block 600, which may

represent the research assistant function being called by another function such as the speech recognition and analysis function illustrated in FIG. 4. At block 610, the function receives search results, i.e. information retrieved from one or more databases, relating to possible actions to solve a problem. In another example, search and retrieval could be performed by a resolution assistant function. In this example, at block 620, the function ranks the possible actions, according to predicted benefits, cost, or some other criteria. At block 630, the function initiates at least one action to solve a problem affecting the customer. This may involve preparing an e-mail message, preparing an order form, preparing an address label, or routing a telephone call, for example. Next, the function may seek approval for actions before completing them. At block 640, the function outputs suggestions for action to a customer service representative (e.g. via audio output, or by displaying the information on a screen, or printing information on paper). For example, actions predicted to be more useful could be highlighted, listed first, or otherwise presented more prominently than other information. If a customer service representative does not approve a suggested action at decision 650, the "No" branch is taken to block 660. There the list is revised (perhaps the rejected action is deleted or moved to the bottom of the list, for example) before the function returns to block 630. If on the other hand a customer service representative approves a suggested action at decision 650, the "Yes" branch is taken to block 670. There the approved action is completed. This may involve completion of a suggested telephone-call routing to another customer service representative, sending an email message containing an advisory bulletin, or ordering goods such as an

updated product, for example. Completion of the action may require further involvement by some person, or completion of the action may be entirely automated. In any event, from the viewpoint of the customer service representative, some or all tasks are delegated or "off-loaded" during a conversation involving a customer.

Next, if more actions remain on the list of possible actions, at decision 680 the "Yes" branch is taken to block 660. There the list of possible actions is revised (the completed action is deleted) before the function returns to block 630. If on the other hand no more actions remain on the list of possible actions, at decision 680 the "No" branch is taken to block 690, where the function terminates (e.g. at the end of a customer service representative's interaction with a customer).

Those skilled in the art will recognize that blocks in the above-mentioned flow charts could be arranged in a somewhat different order, but still describe the invention. Blocks could be added to the above-mentioned flow charts to describe window-managing details, or optional features; some blocks could be subtracted to show a simplified example.

FIG. 7 illustrates selected operations of another exemplary system, employing a research assistant component or a resolution assistant component. This example includes call-routing. Speech recognition and analysis component 232 is shown receiving customer speech input 710, which may be de-cluttered as described above. (Speech recognition and analysis component 232, key word database 235, and one or more databases 260 have been described

above.) In this example, speech recognition and analysis component 232 performs speech recognition on said speech input to generate a text equivalent, and parses said text to identify a key word. A research assistant component or a resolution assistant component at 720 searches for an occurrence of said key word in a database 260 and retrieves information from database 260. The retrieving is completed during a conversation involving a customer and a customer service representative.

Either the research assistant component or resolution assistant component at 720 may be implemented by itself. The two may be implemented in the same system as separate components, or as a combination component. They may be implemented with well-known software technologies, including a web server, database management software, and a scripting language such as PHP, Perl, or another similar language, for example. Other technology similar to the above-mentioned examples also may be used.

At arrow 730, a research assistant component or resolution assistant component at 720 provides retrieved information to a first customer service representative through interface 750. Interfaces 750 and 770 may for example represent audio output, or displaying the information on a screen, or printing information on paper. As one example, when a customer speaks the name of a certain product, a research assistant component or resolution assistant component at 720 would access a database pertaining to that product. Next, when a customer describes a problem, a research assistant component at 720 would provide relevant information for diagnosing the problem to a first customer service representative, through interface 750. At the same time,

a resolution assistant component at 720 may initiate corrective action, such as ordering a repair package or sending instructions to the customer by email.

5 The output at arrow 730 may include suggested actions such as routing a customer's telephone call to a second customer service representative who is a specialist, for example. At arrow 740, the first customer service representative may signal approval of suggested actions such as call routing. The customer's call may 10 then be routed to a specialist who would continue the conversation with the customer (resulting in further speech input from the customer), and solve the customer's problem. The second representative may be more skilled in the problem area. The two representatives may be at the same location or at different 15 locations. The two representatives may use the same system, or separate systems. A decision to transfer handling to another customer service representative may be based on one or more key words (e.g. a customer name, product name, or problem description), automatically recognized in the customer's speech. 20 Thus the system in this example initiates action, based on a key word, to solve a problem affecting a customer.

25 The system in this example creates or updates a problem management record (PMR 760) to document research or actions to solve a problem affecting the customer. Research assistant component or resolution assistant component 720 provides information via PMR 760 to the second customer service representative (the specialist) through interface 770. This may be accomplished by utilizing a communications network (not shown) 30 such as a local area network or the Internet, for example.

In conclusion, examples have been shown of methods and systems employing computerized speech recognition and key words to improve customer service.

5 One of the preferred implementations of the invention is an application, namely a set of instructions (program code) in a code module which may, for example, be resident in the random access memory of a computer. Until required by the computer, the set of instructions may be stored in another computer memory, for example, in a hard disk drive, or in a removable memory such as an optical disk (for eventual use in a CD ROM) or floppy disk (for eventual use in a floppy disk drive), or downloaded via the Internet or other computer network. Thus, the present invention may be implemented as a computer-readable medium having computer-executable instructions for use in a computer. In addition, although the various methods described are conveniently implemented in a general-purpose computer selectively activated or reconfigured by software, one of ordinary skill in the art would also recognize that such methods may be carried out in hardware, in firmware, or in more specialized apparatus constructed to perform the required method steps.

20 While the invention has been shown and described with reference to particular embodiments thereof, it will be understood by those skilled in the art that the foregoing and other changes in form and detail may be made therein without departing from the spirit and scope of the invention. The appended claims are to encompass within their scope all such changes and modifications as are within the true spirit and scope of this invention. Furthermore, 25 it is to be understood that the invention is solely defined by 30

the appended claims. It will be understood by those with skill in the art that if a specific number of an introduced claim element is intended, such intent will be explicitly recited in the claim, and in the absence of such recitation no such limitation is present. For non-limiting example, as an aid to understanding, the appended claims may contain the introductory phrases "at least one" or "one or more" to introduce claim elements. However, the use of such phrases should not be construed to imply that the introduction of a claim element by indefinite articles such as "a" or "an" limits any particular claim containing such introduced claim element to inventions containing only one such element, even when the same claim includes the introductory phrases "at least one" or "one or more" and indefinite articles such as "a" or "an;" the same holds true for the use in the claims of definite articles.